

A Minimax Distortion View of Differentially Private Query Release

Weina Wang, Lei Ying, and Junshan Zhang

Abstract

We consider the problem of differentially private query release through a synthetic database approach. Departing from the existing approaches that require the query set to be specified in advance, we advocate to devise query-set independent mechanisms, with an ambitious goal of providing accurate answers, while meeting the privacy constraints, for all queries in a general query class. Specifically, a differentially private mechanism is constructed to “encode” rich stochastic structure into the synthetic database, and “customized” companion estimators are then derived to provide accurate answers by making use of all available information, including the mechanism (which is public information) and the query functions. Accordingly, the distortion under the best of this kind of mechanisms at the worst-case query in a general query class, so called the minimax distortion, provides a fundamental characterization of differentially private query release.

For the general class of statistical queries, we prove that with the squared-error distortion measure, the minimax distortion is $O(1/n)$ by deriving asymptotically tight upper and lower bounds in the regime that the database size n goes to infinity. The upper bound is achievable by a mechanism \mathcal{E} and its corresponding companion estimators, which points directly to the feasibility of the proposed approach in large databases. We further evaluate the mechanism \mathcal{E} and the companion estimators through experiments on real datasets from Netflix and Facebook. Experimental results show improvement over the state-of-art MWEM algorithm and verify the scaling behavior $O(1/n)$ of the minimax distortion.

I. INTRODUCTION

It is envisaged that in the forthcoming “big data” era, there will be an abundance of rich data about individuals in many domains, such as healthcare, mobile networks, social networks and web search. While data analysis uncovers scientific and societal insights, it also poses potential “threats” to personal

W. Wang, L. Ying and J. Zhang are with the School of Electrical, Computer and Energy Engineering, Arizona State University, Tempe, AZ 85281 USA (e-mail: weina.wang@asu.edu; lei.ying.2@asu.edu; junshan.zhang@asu.edu).

privacy. It is therefore of great interest to establish a systematic understanding of privacy-preserving data analysis, aiming to provide utility for data analytics while preserving privacy. To rigorously quantify privacy, the celebrated notion of differential privacy, introduced in a line of work [1]–[3], has emerged as an analytical foundation for privacy-preserving data analysis.

Viewing a database as a vector of rows, with each row corresponding to some sensitive record of an individual (e.g., a patient’s medical record), an information releasing mechanism is said to be ϵ -*differentially private* if the change of a single row alters the probability of any output instance by at most an e^ϵ multiplicative factor. By this requirement, the presence of an individual, or the content of the record associated with an individual, cannot be exactly deduced from the released information. Therefore, a differentially private mechanism guarantees that only limited *additional* information about an individual would be leaked.

As is standard, information about a database is acquired through queries. Therefore, a central problem in differential privacy is to privately release outputs that permit accurate answers to be derived for as many as possible queries. This problem has been extensively studied in the differential privacy literature, and many mechanisms for query release have been developed (see, e.g., [1], [4]–[11]). Adopted by much of the existing work, a natural approach is to non-interactively generate a synthetic database, which is a one-shot “sanitization” of the original database consisting of rows that come from the same data universe as the rows of the original database.

In contrast to the interactive counterpart, the non-interactive synthetic database approach allows arbitrary number of queries to be answered without compromising differential privacy. More specifically, queries arrive online in the interactive approach and each query consumes some privacy budget. Therefore, a privacy allocation plan is needed and only a finite number of queries can be answered before the privacy is breached. While in the non-interactive approach, the privacy budget is used all at once for the synthetic database generation, since further processing of the released synthetic database does not consume any privacy budget. As long as the synthetic database is released through a differentially private mechanism, arbitrary number of queries can be answered without compromising differential privacy.

However, although the synthetic database approach allows arbitrary number of queries to be answered without compromising differential privacy, most existing mechanisms for synthetic database release are still confined to a specific query set. Typically, the existing mechanisms [4], [5], [9]–[11] require the query set to be specified beforehand, and the accuracy guarantee becomes worse as the size of the query set increases. There are at least two drawbacks in this approach. First, to specify a query set beforehand, a priori knowledge of the queries of interest is needed, and the queries cannot be chosen adaptively.

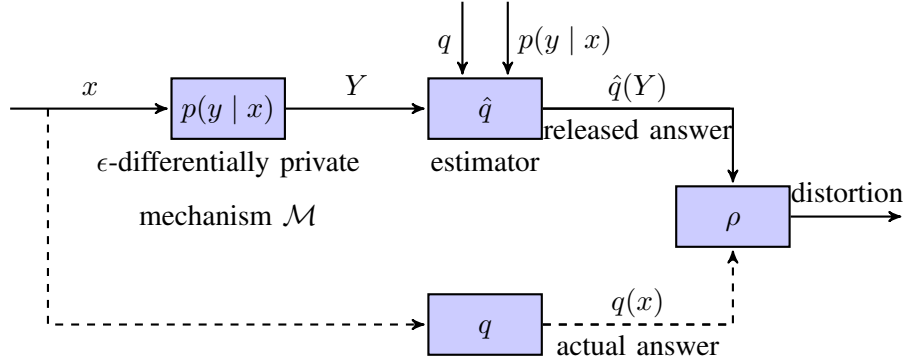


Fig. 1: Road map of our approach for differentially private query release.

Second, to achieve certain accuracy, the size of the query set must be smaller than a threshold. However, as pointed out in [12], in many research settings, it is hard to decide in advance exactly which statistics should be computed. As a consequence, the synthetic database approach would not work well for such scenarios. These drawbacks debilitate the promise that arbitrary number of queries can be answered privately in a non-interactive approach, giving rise to the following question: *is it possible to make the synthetic database releasing mechanism independent of any specific query set while still enabling accurate answers to be derived for all queries in a general query class from the released synthetic database?* If this could be done, the synthetic database approach would be literally “non-interactive,” in the sense that users do not need to interact with the curator during the entire process, whereas users need to submit the query set to the curator beforehand in the existing mechanisms.

In this paper, we give positive answers to the above question for a general class of queries, via taking the following approach. First, a synthetic database is released by a query-set independent differentially private mechanism, aiming at providing accurate answers for all queries in the query class. Then each query is answered by an estimator based on the released synthetic database, rather than directly carried out as if the synthetic database were the actual database. In particular, the mechanism is constructed to “encode” rich stochastic structure into the synthetic database, and the estimator makes use of the structure of the mechanism (which is public information) and the query function. This approach decouples synthetic database generating and query answering. By introducing the flexibility of “customizing” estimators for different queries, it opens the possibility of deriving accurate answers for all queries in a general query class from the same released synthetic database. We use *synthetic database release* to refer to the process of generating an output synthetic database, and *query release* to refer to the entire process including

releasing a synthetic database and deriving answers to queries using estimators.

Based on this approach, we advocate a minimax distortion view of differentially private query release. Consider a database consisting of n rows/entries, each of which takes values from a domain $\mathcal{D} = \{0, 1\}^l$, i.e., they have l binary attributes. The database is then represented by a vector $x \in \mathcal{D}^n$. Consider an ϵ -differentially private mechanism \mathcal{M} for synthetic database release and let $Y = \mathcal{M}(x)$ denote the output. For each query $q: \mathcal{D}^n \rightarrow \mathcal{R}$ in a query class \mathcal{Q} , where \mathcal{R} is some abstract space, an estimator $\hat{q}: \mathcal{D}^n \rightarrow \mathcal{R}$ is used to answer the query based on the synthetic database, and the answer is denoted by $\hat{q}(Y)$, as illustrated in Figure 1. The accuracy of \mathcal{M} for a query $q \in \mathcal{Q}$ is evaluated when an optimal estimator \hat{q}^* is in use, since an optimal estimator fully exploits the available information in the mechanism. To guarantee accuracy for all queries in the query class, the performance of \mathcal{M} is measured by the worst-case distortion among queries in \mathcal{Q} . Then a fundamental characterization of differentially private query release is the following minimax distortion:

$$\mathfrak{D}_\epsilon = \inf_{\substack{\epsilon\text{-differentially} \\ \text{private mechanisms}}} \sup_{q \in \mathcal{Q}, x \in \mathcal{D}^n} \mathbb{E}[\rho(\hat{q}^*(Y), q(x))], \quad (1)$$

where ρ is a distortion measure, \hat{q}^* is the optimal estimator, and Y follows the probability distribution induced by x through the mechanism. This minimax distortion characterizes the best one can get from an ϵ -differentially private synthetic database releasing mechanism for the worst-case query accuracy guarantee, yielding a minimax distortion view of differentially private query release. Our main contributions are summarized as follows.

Contributions

1) We propose a two-phase approach for differentially private query release: First, a synthetic database is released by a query-set independent differentially private mechanism, aiming at providing accurate answers for all queries in a general query class; Then queries are answered by customized estimators. Based on this approach, we advocate a minimax distortion view of differentially private query release, where the minimax distortion \mathfrak{D}_ϵ is defined to be the distortion under the best ϵ -differentially private synthetic database releasing mechanism for the worst-case query in a general query class. Accordingly, the best mechanism allows all queries in a general query class to be answered with a distortion upper bounded by the minimax distortion.

2) For the class of statistical queries (which is a generalization of the class of linear queries in the literature), we consider the minimax distortion \mathfrak{D}_ϵ^S with the squared-error distortion measure, i.e., $\rho(s, t) = (s - t)^2$ for any $s, t \in \mathbb{R}$. We prove that the minimax distortion \mathfrak{D}_ϵ^S is $O(1/n)$ by deriving

asymptotically tight upper and lower bounds in the regime that the database size n goes to infinity, for given data universe dimension l and privacy level ϵ .

The upper bound on \mathfrak{D}_ϵ^S is achieved by a differentially private synthetic database releasing mechanism \mathcal{E} and the companion estimators. The mechanism \mathcal{E} can be viewed as an instance of the exponential mechanism and the randomized response mechanism. It encodes an independence structure into the released synthetic database that is exploited by the companion estimators. Under \mathcal{E} and the estimators, all the statistical queries can be answered with distortion $O(1/n)$, which guarantees reasonable accuracy in large databases. In conclusion, there exists a query-set independent differentially private synthetic database releasing mechanism that permits accurate answers to be derived for all the statistical queries from the released synthetic database.

3) We evaluate the mechanism \mathcal{E} and the companion estimators through a number of experiments. The experimental results on a Netflix dataset for statistical queries show that this approach provides reasonable accuracy for all the tested queries, irrespective of the form of the queries or the number of the tested queries, which improves over the MWEM algorithm. The scaling behavior $O(1/n)$ of the minimax distortion is also verified by the results. The experiment on a Facebook dataset shows that this approach works well for the application of differentially private cut function release for graphs.

Related Work

Differential privacy, introduced in the seminal work [1], [2], has attracted much attention and has emerged as an analytical foundation for privacy-preserving data analysis. Extensive research has been done for both interactive and non-interactive approaches.

Non-interactive approaches have been preferred by data-mining and statistics community. However, some negative results have been found about this approach. Dinur and Nissim [13] showed that noise of magnitude $o(\sqrt{n})$ is blatantly non-private against $n \log^2 n$ random queries, where the queries may involve only a subset of the rows. Dwork et al. [1] considered the statistical difference between two distributions that are induced by two databases that have very different answers to the same query. They showed that for many queries, this statistical difference is small unless the database size is exponential in the dimension of the data universe.

These negative results motivate interactive approaches, where the number of queries was initially limited to a sublinear order of n . Dwork et al. [1] proposed the Laplace mechanism that adds Laplace noise to the real answer of a low sensitivity query. When independent noise is added to different queries, the distortion of each query scaled as $O(|\mathcal{Q}|/n)$. Subsequent work [6], [8] focused on predicate/linear queries

and developed mechanisms that allow exponential number of queries to be answered with distortion $O(\text{polylog}(|\mathcal{Q}|)/n^{1/3})$ and $O((\log(|\mathcal{Q}|))^{1/2}/n^{1/2})$, respectively, where the latter is for (ϵ, δ) -differential privacy.

Non-interactive approach was revisited by Blum, Ligett and Roth [4]. The mechanism proposed in this work guarantees that the distortion for each predicate query in a concept class \mathcal{Q} is upper bounded by $O((\text{VCDIM}(\mathcal{Q}))^{1/3}/n^{1/3})$, where $\text{VCDIM}(\mathcal{Q})$ is the VC-dimension of \mathcal{Q} . A similar distortion bound $O((\log(|\mathcal{Q}|))^{1/3}/n^{1/3})$ was achieved by the work of Hardt, Ligett and McSherry [9] for linear queries. A distortion bound $O((\log(|\mathcal{Q}|))^{1/2}/n^{1/2})$ under (ϵ, δ) -differential privacy was also achieved in this work. In this paper, we consider a more general class of queries, named statistical queries, and aim at providing accurate answers for all queries in this query class. If the absolute-error distortion $\rho = |s - t|$ for any $s, t \in \mathbb{R}$ is used, as the above related work, then the synthetic database releasing mechanism \mathcal{E} and the proposed companion estimators give answers to all the statistical queries with expected distortion $O(1/n^{1/2})$.

Minimax risk is a classical framework in statistics [14] that focuses on estimating parameters of the underlying distribution. Minimax rates were studied under *local privacy*, which is a privacy notion different from differential privacy, by Duchi, Jordan and Wainwright [15], [16]. In contrast, this study does not assume any knowledge of the underlying distribution of the database, and focuses on providing accurate answers to a general class of queries.

Paper Organization

The rest of the paper is organized as follows. In Section II, we describe the model used in this paper. In Section III, we present our minimax distortion view of the differentially private query release. The class of statistical queries is studied in Section IV, and some generalizations are given in Section V. Experimental evaluation of the proposed approach and the application of cut function release for graphs are presented in Section VI. Finally, we conclude our work and discuss future work in Section VII.

Notation: Throughout this paper we use the following basic notation. Denote the set of real numbers by \mathbb{R} , the set of nonnegative real numbers by \mathbb{R}^+ . Let $\overline{\mathbb{R}}^+ = \mathbb{R}^+ \cup \{+\infty\}$. Denote the set of nonnegative integers by \mathbb{N} and denote $[n] = \{1, 2, \dots, n\}$ for $n \in \mathbb{N} \setminus \{0\}$.

II. MODEL

We consider the following model for a database. A database is represented by a vector x of length n , with each entry corresponding to a row of the database and n being the size of the database. Entries of x

are denoted by x_1, x_2, \dots, x_n , and they take values from a domain $\mathcal{D} = \{0, 1\}^l$, i.e., they have l binary attributes. Then $\mathcal{D}^n = (\{0, 1\}^l)^n$ denotes the set of all possible databases. Two databases $x, x' \in \mathcal{D}^n$ are said to be *neighbors* if they differ on exactly one row, and $x \sim x'$ denotes the neighboring relation.

Information about a database is acquired through queries. A *query* is a function $q: \mathcal{D}^n \rightarrow \mathcal{R}$, where \mathcal{R} is some abstract range. Consider a database $x \in \mathcal{D}^n$. The answer $q(x)$ to the query contains information about x ; however, directly releasing $q(x)$ may compromise privacy, necessitating privacy-preserving information releasing mechanisms.

Definition 1. A *mechanism* \mathcal{M} is specified by an *associated mapping* $\mu_{\mathcal{M}}: \mathcal{D}^n \rightarrow \mathcal{P}$, where \mathcal{P} is the set of probability measures on some measurable space $(\mathcal{S}, \mathcal{F})$, called the *range* of the mechanism \mathcal{M} . Taking a database $x \in \mathcal{D}^n$ as the input, the mechanism \mathcal{M} outputs an \mathcal{S} -valued random variable with distribution measure $\mu_{\mathcal{M}}(x)$ on $(\mathcal{S}, \mathcal{F})$.

Definition 2. (Dwork et al. [1], [2]) A mechanism \mathcal{M} is ϵ -*differentially private* for some $\epsilon \in \overline{\mathbb{R}}^+$ if for any pair of neighboring databases $x, x' \in \mathcal{D}^n$, and any measurable $\mathcal{K} \in \mathcal{F}$,

$$\mathbb{P}\{\mathcal{M}(x) \in \mathcal{K}\} \leq e^\epsilon \mathbb{P}\{\mathcal{M}(x') \in \mathcal{K}\}. \quad (2)$$

Intuitively, differential privacy requires certain indistinguishability between the distributions induced by neighboring databases. The smaller ϵ is, the more indistinguishability is required, and hence the better privacy is. We call the parameter ϵ the *level of differential privacy*. Note that the differential privacy property of a mechanism is fully characterized by its associated mapping.

We consider differentially private mechanisms for non-interactive synthetic database release. Specifically, let $\wp(\mathcal{D}^n)$ denote the power set of \mathcal{D}^n . Then we consider differentially private mechanisms with range $(\mathcal{D}^n, \wp(\mathcal{D}^n))$. Let \mathcal{M} be such a mechanism and $x \in \mathcal{D}^n$ be a database. Then the output $Y = \mathcal{M}(x)$ is a \mathcal{D}^n -valued random variable that represents the released synthetic database. Many mechanisms for synthetic database release have been developed (see, e.g., [4], [5], [9]–[11]), where a query q is typically answered by $q(Y)$, i.e., a query is answered as if the synthetic database were the actual database. These mechanisms require the query set to be specified in advance and the accuracy guarantee depends on the size of the query set.

In this paper, we explore the following approach. First, a synthetic database is released using a query-set independent differentially private mechanism, and then queries are answered by customized estimators. For each query q in a query class \mathcal{Q} , an estimator $\hat{q}: \mathcal{D}^n \rightarrow \mathcal{R}$ is used to answer the query based on the synthetic database, and thus the answer is denoted by $\hat{q}(Y)$. To achieve good accuracy, the estimator

\hat{q} should be designed according to the mechanism \mathcal{M} and the query q , making use of all the available information. The distortion between the actual answer $q(x)$ and the released answer $\hat{q}(Y)$ is measured by a distortion measure ρ on the range of the query q . This approach is illustrated in Figure 1, where the mechanism is represented by the probability distribution $p(\cdot | x)$ of Y for each input database x , and \hat{q} has $p(\cdot | x)$ and q as inputs to indicate the design dependence.

Note that in this non-interactive approach, as long as the mechanism \mathcal{M} is ϵ -differentially private, the whole query release process is ϵ -differentially private, i.e., arbitrary number of queries can be answered and any estimator can be used, with the level of differential privacy still preserved.

III. MINIMAX DISTORTION

The proposed approach aims at privately releasing a synthetic database that permits accurate answers to be derived for all queries in a query class. Therefore, a natural fundamental characterization of differentially private query release is the following minimax distortion: the distortion under the best differentially private synthetic database releasing mechanism (the “min” part) for the worst-case query in the query class (the “max” part).

For each query $q: \mathcal{D}^n \rightarrow \mathcal{R}$, let $\rho: \mathcal{R} \times \mathcal{R} \rightarrow \mathbb{R}^+$ be a distortion measure on the space \mathcal{R} . For the sake of fair comparison, we assume that q is normalized, i.e.,

$$\max_{x, x' \in \mathcal{D}^n} \rho(q(x), q(x')) = 1, \quad (3)$$

which rules out trivial queries that map all possible databases to a constant. For each query q , to guarantee that the released answers have “physical meanings,” we consider the estimators such that the answers released by them correspond to possible answers to the query q on real databases, i.e., the estimators in $\hat{\mathcal{Q}}_q = \{\hat{q} \rightarrow \mathcal{R} \mid \hat{q}(\mathcal{D}^n) \subseteq q(\mathcal{D}^n)\}$, which we call *proper estimators*. Consider an ϵ -differentially private mechanism \mathcal{M} and an estimator $\hat{q} \in \hat{\mathcal{Q}}_q$ for the query q , the distortion of the answer is defined by the following worst-case distortion among all possible databases:

$$\sup_{x \in \mathcal{D}^n} \mathbb{E}_{Y \sim \mu_{\mathcal{M}}(x)} [\rho(\hat{q}(Y), q(x))], \quad (4)$$

where the subscript $Y \sim \mu_{\mathcal{M}}(x)$ indicates that Y follows the distribution $\mu_{\mathcal{M}}(x)$, and the expectation is taken over all the randomness.

To minimize distortion, an estimator should be designed according to the mechanism \mathcal{M} and the query q , making use of all the available information. Therefore an optimal estimator \hat{q}^* is given by

$$\hat{q}^* \in \arg \inf_{\hat{q} \in \hat{\mathcal{Q}}_q} \sup_{x \in \mathcal{D}^n} \mathbb{E}_{Y \sim \mu_{\mathcal{M}}(x)} [\rho(\hat{q}(Y), q(x))]. \quad (5)$$

Note that the set $\hat{\mathcal{Q}}_q$ contains only a finite number of estimators since it consists of mappings from \mathcal{D}^n to $q(\mathcal{D}^n)$, which are both finite sets, indicating that the infimum in (5) can be attained. Since the information in a mechanism is fully exploited only when an optimal estimator is in use, the accuracy of an ϵ -differentially private mechanism \mathcal{M} for a query q is evaluated with an optimal estimator \hat{q}^* , i.e., by the distortion

$$\sup_{x \in \mathcal{D}^n} \mathbb{E}_{Y \sim \mu_{\mathcal{M}}(x)} [\rho(\hat{q}^*(Y), q(x))]. \quad (6)$$

The synthetic database released by \mathcal{M} is expected to answer all queries in a query class \mathcal{Q} . To guarantee accuracy for all queries in \mathcal{Q} , the performance of \mathcal{M} is measured by the worst-case distortion among all queries in \mathcal{Q} , i.e., by

$$\sup_{q \in \mathcal{Q}} \left\{ \sup_{x \in \mathcal{D}^n} \mathbb{E}_{Y \sim \mu_{\mathcal{M}}(x)} [\rho(\hat{q}^*(Y), q(x)) \right\}. \quad (7)$$

Let \mathcal{U}_ϵ be the set of mappings associated with ϵ -differentially private mechanisms. Then we define the *minimax distortion* as

$$\mathfrak{D}_\epsilon = \inf_{\mu_{\mathcal{M}} \in \mathcal{U}_\epsilon} \sup_{q \in \mathcal{Q}} \left\{ \sup_{x \in \mathcal{D}^n} \mathbb{E}_{Y \sim \mu_{\mathcal{M}}(x)} [\rho(\hat{q}^*(Y), q(x)) \right\}. \quad (8)$$

The minimax distortion is a fundamental characterization of ϵ -differentially private query release since it characterizes the best one can get from an ϵ -differentially private synthetic database releasing mechanism for the worst-case query accuracy guarantee. In what follows we will study differentially private query release from this minimax distortion view, and derive upper and lower bounds on the minimax distortion accordingly.

IV. STATISTICAL QUERIES

In this section, we consider differentially private query release for the class of statistical queries, which is a much larger class than the class of linear queries in the literature.

Definition 3. A *statistical query* $q_\varphi: \mathcal{D}^n \rightarrow \mathbb{R}$ is specified by a sequence of functions

$$\varphi = (\varphi_i: \mathcal{D} \rightarrow \mathbb{R}, i = 1, 2, \dots), \quad (9)$$

where each φ_i is a function of the i th row of the database, which we call a *row function*, and there is no constraint on its form except boundedness. Let $a_i = \min_{v \in \mathcal{D}} \varphi_i(v)$, $b_i = \max_{v \in \mathcal{D}} \varphi_i(v)$ and $c_i = b_i - a_i$. Assume that for any $i \in [n]$, $a \leq a_i < b_i \leq b$ and $c_i \geq c$ for some $a, b, c \in \mathbb{R}$ with $c > 0$. Then q_φ is defined by

$$q_\varphi(x) = \frac{1}{\sum_{i=1}^n c_i} \sum_{i=1}^n \varphi_i(x_i), \quad (10)$$

where x_1, \dots, x_n are the rows of the database x .

Note that the above definition of statistical query is a generalization of the so called *linear query* (and its special form *predicate/counting query*) in the literature [4], [6], [8]–[10], [17]–[19], since a linear query can be written as a statistical query with identical row functions for all the rows. Linear queries can be answered as long as the histogram of a database is known. However, histograms are often not sufficient for answering statistical queries, making the approaches that privately release histograms not applicable for statistical queries.

Denote the class of statistical queries by \mathcal{Q}^S and let $\rho: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}^+$ be the squared-error distortion, i.e., $\rho(s, t) = (s - t)^2$ for any $s, t \in \mathbb{R}$. Then the minimax distortion for statistical queries can be written as

$$\mathfrak{D}_\epsilon^S = \inf_{\mu, \mathcal{M} \in \mathcal{U}_\epsilon} \sup_{q_\varphi \in \mathcal{Q}^S, x \in \mathcal{D}^n} \mathbb{E}_{Y \sim \mu, \mathcal{M}(x)} [|\hat{q}_\varphi^*(Y) - q_\varphi(x)|^2]. \quad (11)$$

Theorem 1. *The minimax distortion for statistical queries satisfies the following bounds:*

$$\frac{(1 - \Phi(1))^2}{2^{l+4}(1 + \frac{e^\epsilon}{2^l - 1})^3} \frac{1}{n} + o\left(\frac{1}{n}\right) \leq \mathfrak{D}_\epsilon^S \leq \frac{4(b - a)^2(1 + (2^l - 1)e^{-\epsilon})^2}{c^2(1 - e^{-\epsilon})^2} \frac{1}{n}, \quad (12)$$

where Φ is the cumulative distribution function (cdf) of the standard Gaussian distribution, and a, b, c are the constants in Definition 3.

The upper bound in this theorem is given by the performance of an ϵ -differentially private synthetic database releasing mechanism \mathcal{E} and the companion estimators, which are presented in Section IV-A. The lower bound in this theorem is derived by bounding the average distortion. The minimax distortion is defined for the worst-case distortion over statistical queries and databases. We consider a stochastic model for the queries and the database. Then the average distortion under this model serves as a lower bound on the worst-case distortion. Analyzing the average distortion under the constraint of ϵ -differential privacy as in Section IV-B gives the lower bound.

Consider the asymptotic regime that the database size n goes to infinity for given data universe dimension l and privacy level ϵ . Then the upper bound indicates that there exist query-set independent differentially private synthetic database releasing mechanisms and estimators such that all the statistical queries can be answered with distortion $O(1/n)$. Further, the lower bound and the upper bound are of the same order in terms of database size, which shows that these bounds are asymptotically tight in the considered regime. We derive these bounds in the following subsections.

Remark. We caution that when the privacy level ϵ also scales, the upper and lower bounds given here may not meet. For example, let $\epsilon = n^{-\beta}$ for some $\beta > 0$ and consider the joint asymptotic regime on

the 2-dimensional $(n, \frac{1}{\epsilon})$ -plane. In this case, the upper and lower bounds differ by a factor of the order of $n^{2\beta}$.

A. Upper Bound on the Minimax Distortion

In this subsection, we consider a specific ϵ -differentially private mechanism \mathcal{E} and develop the estimators companioned with it for statistical queries. Since the minimax distortion for statistical queries can be written as

$$\mathfrak{D}_\epsilon^S = \inf_{\mu_{\mathcal{M}} \in \mathcal{U}_\epsilon} \sup_{q_\varphi \in \mathcal{Q}^S} \left\{ \inf_{\hat{q}_\varphi \in \hat{\mathcal{Q}}_{q_\varphi}} \sup_{x \in \mathcal{D}^n} \mathbb{E}_{Y \sim \mu_{\mathcal{M}}(x)} [|\hat{q}_\varphi(Y) - q_\varphi(x)|^2] \right\}, \quad (13)$$

the distortion under the mechanism \mathcal{E} and the developed estimators serves as an upper bound on \mathfrak{D}_ϵ^S , which proves the upper bound in Theorem 1. Since we only consider the mechanism \mathcal{E} in this subsection, we drop the subscript $Y \sim \mu_{\mathcal{M}}(x)$ from expectations for conciseness.

Consider a synthetic database releasing mechanism \mathcal{E} with associated mapping $\mu_{\mathcal{E}}$. For each database $x \in \mathcal{D}^n$, since the output $\mathcal{E}(x)$ has a discrete alphabet \mathcal{D}^n , we use the pmf $p_{\mathcal{E}(x)}$ to represent the distribution measure $\mu_{\mathcal{E}}(x)$. Then let the mechanism \mathcal{E} be specified by

$$p_{\mathcal{E}(x)}(y) = \frac{e^{-\epsilon d(x,y)}}{(1 + (2^l - 1)e^{-\epsilon})^n}, \quad x, y \in \mathcal{D}^n, \quad (14)$$

where $\epsilon \in \mathbb{R}^+$ and d is the Hamming distance on \mathcal{D}^n . By the form of $p_{\mathcal{E}(x)}$, this mechanism can be cast as an instance of the exponential mechanism with score function $-d$ [20].

Let Y denote $\mathcal{E}(x)$ for conciseness when it is clear from the context that x is the underlying database. Then the pmf $p_{\mathcal{E}(x)}$ can be written as

$$p_Y(y) = \prod_{i=1}^n \frac{e^{-\epsilon \delta(x_i, y_i)}}{1 + (2^l - 1)e^{-\epsilon}}, \quad y \in \mathcal{D}^n, \quad (15)$$

where $\delta(x_i, y_i) = 0$ if $x_i = y_i$ and $\delta(x_i, y_i) = 1$ otherwise. Let Y_i denote the i th row of Y . Due to the product form above, the entries $\{Y_i, i \in [n]\}$ are independent and each entry Y_i has the following pmf

$$p_{Y_i}(y_i) = \frac{e^{-\epsilon \delta(x_i, y_i)}}{1 + (2^l - 1)e^{-\epsilon}}, \quad y_i \in \mathcal{D}. \quad (16)$$

Therefore this mechanism can also be viewed as a randomized response scheme, where each individual's data is perturbed independently and then the perturbed data is released. Note that the mechanism \mathcal{E} can be implemented distributedly due to the independence.

The differential privacy property of the mechanism \mathcal{E} is given in the following lemma. The proof is standard and thus we omit it here due to space limit.

Lemma 1. *The mechanism \mathcal{E} is ϵ -differentially private.*

Next we present the estimators companioned with the mechanism \mathcal{E} for the class of statistical queries. Let $g(\epsilon) = 1 + (2^l - 1)e^{-\epsilon}$. For each $q_\varphi \in \mathcal{Q}^S$, consider the estimator $\hat{q}_\varphi^u: \mathcal{D}^n \rightarrow \mathcal{R}$ defined by

$$\hat{q}_\varphi^u(y) = \frac{g(\epsilon)}{1 - e^{-\epsilon}} q_\varphi(y) - \frac{e^{-\epsilon}}{1 - e^{-\epsilon}} C_\varphi, \quad (17)$$

where

$$C_\varphi = \frac{1}{\sum_{i=1}^n c_i} \sum_{i=1}^n \sum_{v \in \mathcal{D}} \varphi_i(v). \quad (18)$$

Lemma 2. *Under the mechanism \mathcal{E} , the estimator \hat{q}_φ^u is unbiased, i.e., for any database $x \in \mathcal{D}^n$,*

$$\mathbb{E}[\hat{q}_\varphi^u(Y)] = q(x), \quad (19)$$

and the distortion of \hat{q}_φ^u satisfies the following upper bound:

$$\sup_{x \in \mathcal{D}^n} \mathbb{E}[|\hat{q}_\varphi^u(Y) - q_\varphi(x)|^2] \leq \frac{(b - a)^2 (1 + (2^l - 1)e^{-\epsilon})^2}{c^2 (1 - e^{-\epsilon})^2} \frac{1}{n}, \quad (20)$$

where a, b, c are the constants in Definition 3.

The proof of this lemma is given in Appendix A. The intuition is that the mechanism \mathcal{E} perturbs each row of the underlying database independently, which encodes an independence structure into the released synthetic base, and then the estimator \hat{q}_φ^u exploits this structure. By the law of large numbers (LLN), the aggregate perturbation converges to the expectation, which is a constant determined by the query and thus can be removed in the estimator.

Next we present a proper estimator designed based on \hat{q}_φ^u . The answer given by the estimator \hat{q}_φ^u may not always be consistent with an actual database, in which case $\hat{q}_\varphi^u \notin \hat{\mathcal{Q}}_{q_\varphi}$. Thus we consider the estimator $\hat{q}_\varphi: \mathcal{D}^n \rightarrow \mathcal{R}$ defined by

$$\hat{q}_\varphi(y) \in \arg \min_{r \in q_\varphi(\mathcal{D}^n)} |\hat{q}_\varphi^u(y) - r|, \quad (21)$$

which quantizes the answer given by \hat{q}_φ^u to the closest value in $q_\varphi(\mathcal{D}^n)$. This quantization guarantees that \hat{q}_φ is a proper estimator, and degrades the performance guarantee only by a factor of 4 as shown in the following lemma, the proof of which is given in Appendix B.

Lemma 3. *Under the mechanism \mathcal{E} , the distortion of the estimator \hat{q}_φ satisfies the following upper bound:*

$$\sup_{x \in \mathcal{D}^n} \mathbb{E}[|\hat{q}_\varphi(Y) - q_\varphi(x)|^2] \leq \frac{4(b - a)^2 (1 + (2^l - 1)e^{-\epsilon})^2}{c^2 (1 - e^{-\epsilon})^2} \frac{1}{n}, \quad (22)$$

where a, b, c are the constants in Definition 3.

Consider the asymptotic regime that the database size n goes to infinity for given data universe dimension l and privacy level ϵ . By the upper bounds (20) and (22), the estimators \hat{q}_φ^u and \hat{q}_φ answer all the statistical queries with distortion $O(1/n)$ based on the synthetic database released by the mechanism \mathcal{E} . Therefore all the statistical queries can be answered with reasonable accuracy guarantee in large databases.

Compared with existing approaches, the synthetic database releasing mechanism \mathcal{E} does not require a priori knowledge of the queries of interest, and instead of answering query q_φ by $q_\varphi(Y)$, the estimators \hat{q}_φ^u and \hat{q}_φ make more use of the stochastic structure in Y encoded by the mechanism \mathcal{E} .

Remark. Under the absolute-error distortion defined by $\rho(s, t) = |s - t|$, for any $s, t \in \mathbb{R}$, the distortion upper bounds for the estimators \hat{q}_φ^u and \hat{q}_φ become

$$\begin{aligned} \sup_{x \in \mathcal{D}^n} \mathbb{E}[|\hat{q}_\varphi^u(Y) - q_\varphi(x)|] &\leq \frac{(b-a)(1 + (2^l - 1)e^{-\epsilon})}{c(1 - e^{-\epsilon})} \frac{1}{\sqrt{n}} \\ \sup_{x \in \mathcal{D}^n} \mathbb{E}[|\hat{q}_\varphi(Y) - q_\varphi(x)|] &\leq \frac{2(b-a)(1 + (2^l - 1)e^{-\epsilon})}{c(1 - e^{-\epsilon})} \frac{1}{\sqrt{n}} \end{aligned}$$

since by Jensen's inequality $(\mathbb{E}[|X|])^2 \leq \mathbb{E}[|X|^2]$ for any random variable X .

Remark. By the form of the estimator \hat{q}_φ^u in (17), the value $C_\varphi = \frac{1}{\sum_{i=1}^n c_i} \sum_{i=1}^n \sum_{v \in \mathcal{D}} \varphi_i(v)$ is needed to answer the query q_φ . In many cases, this value can be easily obtained rather than exhaustive calculation. In such case, the computation in \hat{q}_φ^u is very efficient. Take the following predicate query for an example. Recall that any $v \in \mathcal{D} = \{0, 1\}^l$ is a binary vector $v = (v_1, \dots, v_l)$ of length l . Consider the predicate function $s(v) = v_{j_1} \cdot v_{j_2} \cdot \dots \cdot v_{j_k}$ for some $\{j_1, \dots, j_k\}$ with $1 \leq k \leq l$, which counts the fraction of rows in the database that have value 1 for attributes j_1, \dots, j_k . This predicate query is a statistical query q_φ with $\varphi_i = s$ for any $i \in [n]$. The value C_φ for this query is $C_\varphi = 2^{l-k}$, which can be obtained by simple analysis.

Remark. The estimator \hat{q}_φ^u is more computationally efficient than the estimator \hat{q}_φ since it does not need to find the value closest to $\hat{q}_\varphi^u(Y)$ in $q_\varphi(\mathcal{D}^n)$. Therefore when we are not constricted to the estimators in $\hat{\mathcal{Q}}_{q_\varphi}$, it is more desirable to use the estimator \hat{q}_φ^u from an implementation perspective.

B. Lower Bound on the Minimax Distortion

Consider any ϵ -differentially private mechanism \mathcal{M} . For any query $q_\varphi \in \mathcal{Q}^S$, the form of the optimal estimator depends on q_φ . Therefore with slight abuse of notation, we denote the optimal estimator by the function $\hat{q}^*: \mathcal{D}^n \times \mathcal{Q}^S \rightarrow \mathbb{R}$ and the answer by $\hat{q}^*(Y, q_\varphi)$, where Y is the synthetic database released

by the mechanism \mathcal{M} . Then our goal is to derive a lower bound on the following worst-case distortion:

$$\sup_{q_\varphi \in \mathcal{Q}^S, x \in \mathcal{D}^n} \mathbb{E}_{Y \sim \mu_{\mathcal{M}}(x)} [|\hat{q}^*(Y, q_\varphi) - q_\varphi(x)|^2]. \quad (23)$$

Consider such a type of queries, each of which is specified by an element $z \in \mathcal{D}^n$ and defined by

$$q_z(x) = \frac{1}{n} d(x, z), \quad x \in \mathcal{D}^n,$$

where d is the Hamming distance on \mathcal{D}^n . For any $v, v' \in \mathcal{D}$, let $\delta(v, v') = 0$ if $v = v'$ and $\delta(v, v') = 1$ otherwise. Then the query q_z can be written as

$$q_z(x) = \frac{1}{n} \sum_{i=1}^n \delta(x_i, z_i),$$

from which we can see that the query q_z is a statistical query. Let

$$\mathcal{Q}^Z = \left\{ q_z : \mathcal{D}^n \rightarrow \mathbb{R} \mid q_z(x) = \frac{1}{n} d(x, z), z \in \mathcal{D}^n \right\}. \quad (24)$$

Then $\mathcal{Q}^Z \subseteq \mathcal{Q}^S$, and therefore

$$\begin{aligned} & \sup_{q_\varphi \in \mathcal{Q}^S, x \in \mathcal{D}^n} \mathbb{E}_{Y \sim \mu_{\mathcal{M}}(x)} [|\hat{q}^*(Y, q_\varphi) - q_\varphi(x)|^2] \\ & \geq \sup_{q_z \in \mathcal{Q}^Z, x \in \mathcal{D}^n} \mathbb{E}_{Y \sim \mu_{\mathcal{M}}(x)} [|\hat{q}^*(Y, q_z) - q_z(x)|^2]. \end{aligned}$$

To derive a lower bound on the above supremum, consider \mathcal{D}^n -valued random variables X, Y, Z with the following distributions. The random variable X follows a uniform distribution, i.e., the probability mass function (pmf) $p_X(x) = \frac{1}{2^{nl}}$ for any $x \in \mathcal{D}^n$. Given $X = x$, the conditional pmf of Y is specified by the distribution measure $\mu_{\mathcal{M}}(x)$, i.e., $p_{Y|X}(y | x) = \mathbb{P}\{\mathcal{M}(x) = y\}$ for any $y \in \mathcal{D}^n$. The random variable Z is independent of X and Y , and it also follows a uniform distribution, i.e., the pmf $p_Z(z) = \frac{1}{2^{nl}}$ for any $z \in \mathcal{D}^n$.

Consider the query q_Z , which is the query in \mathcal{Q}^Z specified by Z . Then q_Z is a query chosen from \mathcal{Q}^Z uniformly at random. Due to the independence between Z and (X, Y) , given any $X = x$ and $Z = z$, the conditional pmf $p_{Y|X,Z}(y | x, z) = p_{Y|X}(y | x)$, which corresponds to $\mu_{\mathcal{M}}(x)$. Therefore

$$\begin{aligned} & \sup_{q_z \in \mathcal{Q}^Z, x \in \mathcal{D}^n} \mathbb{E}_{Y \sim \mu_{\mathcal{M}}(x)} [|\hat{q}^*(Y, q_z) - q_z(x)|^2] \\ & = \sup_{q_z \in \mathcal{Q}^Z, x \in \mathcal{D}^n} \mathbb{E} [|\hat{q}^*(Y, q_Z) - q_Z(X)|^2 \mid X = x, Z = z] \\ & \geq \sum_{z \in \mathcal{D}^n, x \in \mathcal{D}^n} \mathbb{E} [|\hat{q}^*(Y, q_Z) - q_Z(X)|^2 \mid X = x, Z = z] p_X(x) p_Z(z) \\ & = \mathbb{E} [|\hat{q}^*(Y, q_Z) - q_Z(X)|^2]. \end{aligned}$$

Note that we construct the random variables X and Z only for the proof. Our result in Theorem 1 does not assume any stochastic model for the database or the query. Note that $\hat{q}^*(Y, q_Z)$ is a function of Y and Z . Since the conditional expectation is precisely the minimum mean square estimator [21], we have

$$\begin{aligned} & \mathbb{E}[|\hat{q}^*(Y, q_Z) - q_Z(X)|^2] \\ & \geq \mathbb{E}[|\mathbb{E}[q_Z(X) | Y, Z] - q_Z(X)|^2] \end{aligned} \quad (25)$$

$$= \frac{1}{n^2} \mathbb{E}[|\mathbb{E}[d(X, Z) | Y, Z] - d(X, Z)|^2]. \quad (26)$$

Recall that the conditional pmf $p_{Y|X}(\cdot | x)$ is specified by the distribution measure $\mu_{\mathcal{M}}(x)$. Then since the mechanism \mathcal{M} is ϵ -differentially private, for any neighboring $x, x' \in \mathcal{D}^n$ and any $y \in \mathcal{D}^n$,

$$p_{Y|X}(y | x) \leq e^\epsilon p_{Y|X}(y | x').$$

This inequality is needed in the proof of the following lemma, which gives a lower bound on the expectation in (26).

Lemma 4. *There exists a constant C such that*

$$\begin{aligned} & \mathbb{E}[|\mathbb{E}[d(X, Z) | Y, Z] - d(X, Z)|^2] \\ & \geq \frac{1}{4} \left((1 - \Phi(1)) \sigma \gamma^{\frac{3}{2}} \sqrt{n} - \frac{C \rho \gamma}{\sigma^3} \right)^2, \end{aligned} \quad (27)$$

where Φ is the cdf of the standard Gaussian distribution,

$$\gamma = \frac{1}{2(1 + \frac{e^\epsilon}{2^{l-1}})}, \quad \sigma^2 = \frac{1}{2^{l-1}}, \quad \rho = \frac{1}{2^{l-1}}. \quad (28)$$

The proof is presented in Appendix C. By this lemma, for any ϵ -differentially private mechanism \mathcal{M} , the distortion is lower bounded as

$$\begin{aligned} & \sup_{q_\varphi \in \mathcal{Q}^S, x \in \mathcal{D}^n} \mathbb{E}_{Y \sim \mu_{\mathcal{M}}(x)} [|\hat{q}^*(Y, q_\varphi) - q_\varphi(x)|^2] \\ & \geq \frac{(1 - \Phi(1))^2}{2^{l+4} (1 + \frac{e^\epsilon}{2^{l-1}})^3} \frac{1}{n} + o\left(\frac{1}{n}\right), \end{aligned} \quad (29)$$

which further implies the lower bound in Theorem 1.

V. GENERALIZATION

In this section, we consider a generalization on the discrete database model and analyze the corresponding minimax distortion.

A. Continuous Data Universe

Consider databases with data universe \mathcal{D} being an interval in the l dimensional real coordinate space \mathbb{R}^l . We assume that l is a constant, so we present the case that $l = 1$ and $\mathcal{D} = [0, 1]$ for clarity. Consider the class of statistical queries with L -Lipschitz row functions, i.e., the query class

$$\mathcal{Q}_L^S = \{q_\varphi \in \mathcal{Q}^S \mid |\varphi_i(u) - \varphi_i(v)| \leq L|u - v|, \text{ for any } u, v \in [0, 1] \text{ and any } i = 1, 2, \dots\}. \quad (30)$$

Then the minimax distortion can be written as

$$\mathfrak{D}_{\epsilon, L}^S = \inf_{\mu_{\mathcal{M}} \in \mathcal{U}_\epsilon} \sup_{q_\varphi \in \mathcal{Q}_L^S, x \in \mathcal{D}^n} \mathbb{E}_{Y \sim \mu_{\mathcal{M}}(x)} [|\hat{q}_\varphi^*(Y) - q_\varphi(x)|^2].$$

We note that the lower bound in Theorem 1 still holds for continuous data universe since $\{0, 1\}^n \subseteq [0, 1]^n$. To obtain an upper bound, we consider the following approach for a database $x \in [0, 1]^n$: first each row of x is discretized into k bits; then the mechanism \mathcal{E} and the companion estimator \hat{q}_φ are used for the discretized database. Denote the discretized database by \hat{x} . Then $\hat{x} \in \{0, \frac{1}{2^k}, \dots, \frac{2^k-1}{2^k}\}^n$. By the discretization precision, $|x_i - \hat{x}_i| \leq \frac{1}{2^k}$ for $i = 1, 2, \dots, n$. Thus for any $q_\varphi \in \mathcal{Q}_L^S$,

$$|q_\varphi(x) - q_\varphi(\hat{x})| \leq \frac{1}{\sum_{i=1}^n c_i} \sum_{i=1}^n |\varphi_i(x_i) - \varphi_i(\hat{x}_i)| \leq \frac{L}{c2^k}.$$

By Lemma 3,

$$\mathbb{E}[|\hat{q}_\varphi(Y) - q_\varphi(\hat{x})|] \leq \frac{2(b-a)(1 + (2^k-1)e^{-\epsilon})}{c(1 - e^{-\epsilon})} \frac{1}{\sqrt{n}},$$

where we omit the subscript $Y \sim p_{\mathcal{E}(x)}$ of the expectation for conciseness. Then

$$\begin{aligned} & \mathbb{E}[|\hat{q}_\varphi(Y) - q_\varphi(x)|^2] \\ & \leq \mathbb{E}[(|q_\varphi(x) - q_\varphi(\hat{x})| + |\hat{q}_\varphi(Y) - q_\varphi(\hat{x})|)^2] \\ & \leq |q_\varphi(x) - q_\varphi(\hat{x})|^2 + |q_\varphi(x) - q_\varphi(\hat{x})| \cdot \mathbb{E}[|\hat{q}_\varphi(Y) - q_\varphi(\hat{x})|] \\ & \quad + (\mathbb{E}[|\hat{q}_\varphi(Y) - q_\varphi(\hat{x})|])^2 \\ & \leq \frac{L^2}{c^2 2^{2k}} + \frac{2(b-a)(1 + (2^k-1)e^{-\epsilon})L}{c^2(1 - e^{-\epsilon})2^k} \frac{1}{\sqrt{n}} \\ & \quad + \frac{4(b-a)^2(1 + (2^k-1)e^{-\epsilon})^2}{c^2(1 - e^{-\epsilon})^2} \frac{1}{n}. \end{aligned}$$

Let $2^{2k} = \sqrt{n}$. We obtain

$$\mathbb{E}[|\hat{q}_\varphi(Y) - q_\varphi(x)|^2] \leq \left(\frac{L^2}{c^2} + \frac{4(b-a)^2 e^{-2\epsilon}}{c^2(1 - e^{-\epsilon})^2} \right) \frac{1}{\sqrt{n}} + o\left(\frac{1}{\sqrt{n}}\right),$$

which gives an upper bound on $\mathfrak{D}_{\epsilon, L}^S$.

Proposition 1. *With continuous data universe $\mathcal{D} = [0, 1]$, the minimax distortion for statistical queries with L -Lipschitz row functions satisfies the following bounds:*

$$\frac{(1 - \Phi(1))^2}{2^5(1 + e^\epsilon)^3} \frac{1}{n} + o\left(\frac{1}{n}\right) \leq \mathfrak{D}_{\epsilon, L}^S \leq \left(\frac{L^2}{c^2} + \frac{4(b-a)^2 e^{-2\epsilon}}{c^2(1 - e^{-\epsilon})^2}\right) \frac{1}{\sqrt{n}} + o\left(\frac{1}{\sqrt{n}}\right). \quad (31)$$

Remark. For a continuous data universe, the optimal estimator \hat{q}_φ^* may not be attainable. In this case, we need to express the minimax distortion $\mathfrak{D}_{\epsilon, L}^S$ in the same form as (13). However, this does not change the arguments for the lower and upper bounds.

VI. EXPERIMENTAL EVALUATION AND APPLICATION

In this section, we first evaluate the mechanism \mathcal{E} in (14) when companioned with the estimator \hat{q}_φ^u in (17) through experiments on a Netflix dataset [22] for statistical queries. During the experiments, we compare our approach with the MWEM algorithm (a combination of the Exponential Mechanism with the Multiplicative Weights update rules) [9]. The main conclusion from the experimental results is that the proposed approach provides reasonable accuracy for all the tested queries, irrespective of the form of the queries or the number of the tested queries, which improves over the MWEM algorithm. The scaling behavior $O(1/n)$ of the minimax distortion as the database size n goes to infinity is also verified by the experimental results.

We next consider the application of differentially private cut function release for graphs and derive an upper bound on the minimax distortion for this application. We evaluate our approach through experiments on a Facebook dataset [23]. The experimental results verify the theoretical upper bound and show that the proposed approach works well for this application.

A. Evaluation for Statistical Queries

In this subsection, we conduct experiments on the Netflix dataset for statistical queries. The Netflix dataset consists of movie ratings from users, with each rating on a scale from 1 to 5 (integral) stars. We treat each rating as a row and model the dataset as a database. To obtain databases with different sizes, we take subsets from the dataset.

The experimental evaluation in this subsection has three focuses: (1) the separation between statistical queries and linear queries, (2) distortion under varying query set size, and (3) scaling behavior of the distortion under varying database size.

1) *Statistical Queries vs. Linear Queries*: The class of statistical queries is much larger than the class of linear queries since a statistical query allows different row functions, whereas a linear query can only have identical row functions. For the private movie rating release application, it is possible to encounter queries that perform different functions on different rows, since different movies or users may belong to different groups and have different weights in a query. We call the number of distinct row functions in a statistical query the *heterogeneity* of the query. Consider a statistical query q_φ and the associated row function sequence $\varphi = (\varphi_1, \dots, \varphi_n)$. If the heterogeneity of q_φ equals to 1, then $\varphi_1 = \dots = \varphi_n$, and thus q_φ is a linear query. If the heterogeneity of q_φ is greater than 1, then not all the φ_i 's are equal. For example, during the experiments in this subsection, when the heterogeneity equal to 2, the statistical query performs one row function for the first half of the rows, and performs another row function for the second half.

The mechanism \mathcal{E} and the companion estimator \hat{q}_φ^u is designed for statistical queries. The upper bound (20) on the distortion of the proposed approach holds for any statistical query, and thus holds for any heterogeneity. The MWEM algorithm is designed for linear queries. To evaluate the MWEM algorithm for statistical queries, we adapt it as follows. For each distinct row function in a statistical query, we treat the set of rows associated with this row function as a “sub-database”. Restricted to this sub-database, the statistical query is a linear query, so we can run the MWEM algorithm on the sub-database to generate a synthetic sub-database. Then the answer to the statistical query is obtained by combining the answers at each sub-database. When there are multiple statistical queries in the query set, we need to divide the database into sub-databases such that restricted to a sub-database, any query in the query set is a linear query. In the experiment, we consider statistical queries with same row functions for ratings of the same movie.

We evaluate the proposed approach and the MWEM algorithm on a database of size $n = 162,567$ from the Netflix dataset, consisting of ratings for 128 movies. Each movie has roughly $1000 \sim 2000$ ratings. Statistical queries are generated randomly in the following way. To specify a row function φ_i , the values $\varphi(1), \varphi(2), \dots, \varphi(5)$ are sufficient. We generate i.i.d. random variables X_1, \dots, X_5 with uniform distribution on $[0, 1]$, and divide them by $\max_i X_i - \min_i X_i$ for normalization. Then these values are used to specify a row function. For a statistical query with heterogeneity h , we generate h row functions independently, and assign each row function to rows corresponding to $1/h$ of the movies. During the experiments, we consider heterogeneity varying from 1 to 128. For each heterogeneity h , we generate a set of 200 statistical queries with heterogeneity h independently. We use the absolute-error distortion measure, i.e., $\rho(s, t) = |s - t|$ for any $s, t \in \mathbb{R}$, since both our approach and the MWEM algorithm have




Fig. 2: Distortion under varying heterogeneity. The proposed approach is robust to heterogeneity, whereas the distortion of the MWEM algorithm grows as the heterogeneity increases.

distortion upper bound under this distortion measure. We measure the worst-case distortion among the queries in the query set, and then take an average over 20 independent runs. The differential privacy level is fixed to $\epsilon = 1$.

Figure 2 compares our approach against the MWEM algorithm with varying heterogeneity. The figure shows that the proposed approach gives similar distortions irrespective of the heterogeneity, whereas under the MWEM algorithm, the distortion grows as the heterogeneity increases. This experimental result shows a separation between statistical queries and linear queries: approaches designed for linear queries cannot be directly applied to statistical queries without performance loss.

2) *Query Set Size-Independent Distortion:* Under most existing mechanisms [4], [5], [9]–[11] for synthetic database release, the accuracy guarantee becomes worse as the query set size increases. Under the MWEM algorithm, the worst-case distortion among the queries in a query set is $O((\log(|Q|))^{1/3})$, where $|Q|$ is the query set size. Our approach does not restrict to a specific query set. The distortion upper bound in (20) holds for all the statistical queries. Therefore, under our approach, the worst-case distortion among the queries in a query set will not grow as the query set size increases.

We evaluate the proposed approach and the MWEM algorithm on databases from the Netflix dataset. We randomly generate linear query sets with the size varying from 64 to 1,048,576, using the same method as in the previous experiments. We still use the absolute-error distortion measure. We measure the worst-case distortion among the queries in the query set and among 50 databases, with database sizes roughly within 1000 \sim 2000. Then the worst-case distortion is averaged over 20 independent runs. The differential privacy level is fixed to $\epsilon = 1$.

Figure 3 compares our approach against the MWEM algorithm with varying query set size. The




Fig. 3: Distortion under varying query set size. The worst-case distortion of the proposed approach does not depend on the query set size, whereas the distortion of the MWEM algorithm grows (slowly) as the query set size increases.




Fig. 4: Distortion under varying database size. In the asymptotic regime that the database size n goes to infinity, the upper bound is $\Theta(1/n)$, so the distortion is $O(1/n)$.

figures shows that the proposed approach gives similar worst-case distortion for different query set sizes. However, for the MWEM algorithm, although very slowly, the worst-case distortion grows as the query set size increases. Therefore, to achieve certain accuracy, this growth indicates that the query set size must be smaller than a threshold. This experimental result verifies the dependence of the distortion on the query set size under the MWEM algorithm, and shows the advantage of our approach.

3) *Scaling Behavior*: Consider the asymptotic regime that the database size n goes to infinity for given data universe dimension and differential privacy level. We have proved that the worst-case squared-error distortion of the mechanism \mathcal{E} when companioned with the estimator \hat{q}_φ^u is $O(1/n)$. To verify this theoretical upper bound, we evaluate the proposed approach on databases from the Netflix dataset. The sizes of the databases vary from 14,559 to 232,944. A linear query set of size 200 is randomly generated

in the same way as the previous experiments and used for all the databases. We use the squared-error distortion measure, i.e., $\rho(s, t) = (s - t)^2$ for any $s, t \in \mathbb{R}$. We measure the worst-case distortion among the queries in the query set, and then take an average over 20 independent runs. The differential privacy level is fixed to $\epsilon = 1$. Figure 4 compares the distortion under the proposed approach with the upper bound in (20), which verifies the asymptotic order $O(1/n)$ of the distortion.

B. Differentially Private Cut Function Release for Graphs

Consider the scenario that the given database is a graph, where the presence of individual edges is sensitive information. Such a graph can represent the online social connections between individuals. To release useful information for graph analysis, a well studied approach is to privately release the cut function of the graph [10], [24], [25].

Let the graph be $G = (V, E)$ and $\wp(V)$ denote the power set of V . Then the cut function $f_G: \wp(V) \times \wp(V) \rightarrow [|E|]$ associated with this graph is defined by

$$f_G(S, T) = |\{(i, j) \in E \mid i \in S, j \in T\}|, \quad (32)$$

which is the number of edges crossing the S, T -cut for any disjoint $S, T \subseteq V$.

We use a database x to represent the graph G . Since differential privacy needs to be preserved for edges, each row of x corresponds to a vertex pair $(i, j) \in V \times V$, where $x_{i,j} = 1$ if $(i, j) \in E$, and $x_{i,j} = 0$ otherwise. Here we use (i, j) to index each row of x . Thus the data universe is $\{0, 1\}$ with dimension $l = 1$ and the database size $n = |V|^2$. Two databases x, x' are neighbors if there exists exactly one vertex pair (i, j) such that $x_{i,j} \neq x'_{i,j}$.

For any disjoint $S, T \subseteq V$, we write $f_G(S, T)$ as a function $q_{S,T}$ of x and call it a *cut query*. Consider the absolute-error distortion measure $\rho(s, t) = |s - t|$ for any $s, t \in \mathbb{R}$. Then the minimax distortion for ϵ -differentially private cut function release can be written as

$$\mathfrak{D}_\epsilon^C = \inf_{\mu, \mathcal{M} \in \mathcal{U}_\epsilon} \sup_{\substack{x \in \{0,1\}^n \\ S, T \subseteq V, S \cap T = \emptyset}} \mathbb{E}_{Y \sim \mu, \mathcal{M}(x)} [|\hat{q}_{S,T}^*(Y) - q_{S,T}(x)|].$$

Consider the statistical query defined in Definition 3. Then a cut query $q_{S,T}$ can be viewed as an unnormalized statistical query over the subset $S \times T \subseteq V \times V$ of all the rows. The row function is $\varphi_{i,j}(x_{i,j}) = x_{i,j}$ since

$$q_{S,T}(x) = \sum_{(i,j) \in S \times T} x_{i,j}. \quad (33)$$

Consider the mechanism \mathcal{E} and estimator $\hat{q}_{S,T}: \{0, 1\}^n \rightarrow \mathbb{R}$ defined by

$$\hat{q}_{S,T}(y) = \frac{1 + e^{-\epsilon}}{1 - e^{-\epsilon}} q_{S,T}(y) - \frac{e^{-\epsilon}}{1 - e^{-\epsilon}} |S||T|, \quad (34)$$

$ V $	577	1154	1731	2308	2885	3462	4039
Error	10.4%	11.7%	8.7%	5.3%	4.7%	5.3%	5.4%

TABLE I: Relative error for cut queries.

which is an adapted version of the estimator \hat{q}_φ^u defined in (17) for the query $q_{S,T}$. Let Y denote the released synthetic database $\mathcal{E}(x)$. By similar analysis as in the proof of Lemma 2, the distortion is bounded as

$$\mathbb{E}_{Y \sim \mu_\epsilon(x)} [|\hat{q}_{S,T}(Y) - q_{S,T}(x)|] \leq \frac{1 + e^{-\epsilon}}{1 - e^{-\epsilon}} \sqrt{|S||T|}. \quad (35)$$

For any $S, T \subseteq V$, $|S||T| \leq |V|^2$. Therefore the minimax distortion is upper bounded as

$$\mathfrak{D}_\epsilon^C \leq \frac{1 + e^{-\epsilon}}{1 - e^{-\epsilon}} |V|. \quad (36)$$

1) *Evaluation on the Facebook Dataset:* We evaluate the proposed approach on databases from the Facebook dataset for the application of cut query release. The Facebook dataset is a graph. Each vertex in the graph represents a user, and an edge between two vertices indicates that they are friends.

Consider the asymptotic regime that the number of vertices $|V|$ goes to infinity. We have proved that the absolute-error distortion for any cut query is $O(|V|)$. To verify this theoretical upper bound, we apply our approach on subgraphs of the graph given by the Facebook dataset. The graph consists of 4039 vertices and 88,234 edges. The number of vertices in the considered subgraphs vary from 577 to 4039. For each subgraph, cut queries are generated randomly in the following way. Half of the vertices are uniformly sampled and this vertex set is denoted by S . Then S and $V - S$ specify a cut query. This choice of cut queries results in the largest upper bound on the distortion as shown in (35). We generate a cut query set consisting of 100 cut queries independently. We measure the worst-case absolute-error distortion among the cut queries in the query set, and then take an average over 10 independent runs. The differential privacy level is fixed to $\epsilon = 1$. Figure 5 compares the distortion under the proposed approach with the upper bound in (35), which verifies the asymptotic order $O(|V|)$ of the distortion. The worst-case relative distortion in Table I shows that the accuracy is reasonable for cut queries.

VII. CONCLUSION AND FUTURE WORK

In this paper, we developed a minimax approach for differentially private query release, where query-set independent differentially private synthetic database releasing mechanisms are devised and the companion




Fig. 5: Distortion of cut queries under varying number of nodes in the graph. In the asymptotic regime that the number of nodes $|V|$ goes to infinity, the upper bound is $\Theta(|V|)$, so the distortion is $O(|V|)$.

estimators are designed to provide accurate answers for all queries in a general query class. For the general class of statistical queries, we proved that with the squared-error distortion measure, the minimax distortion \mathfrak{D}_ϵ^S is $O(1/n)$ by deriving asymptotically tight upper and lower bounds in the regime that the database size n goes to infinity. The upper bound was achieved by a differentially private synthetic database releasing mechanism \mathcal{E} and the companion estimators, which indicates that it is feasible to use query-set independent differentially private synthetic database releasing mechanisms while providing accurate answers for all the statistical queries in large databases.

In this work, we have focused on the minimax distortion in the asymptotic regime that database size n grows. It is also of great interest to quantify the scaling laws of the minimax distortion in the joint asymptotic regime in terms of database size n , data universe dimension l and the differential privacy level ϵ . We are currently investigating this issue and aim at designing better differentially private synthetic database releasing mechanisms for large data universe dimension l and finding tighter lower bounds in terms of ϵ .

REFERENCES

- [1] C. Dwork, F. McSherry, K. Nissim, and A. Smith, “Calibrating noise to sensitivity in private data analysis,” in *Proc. Conf. Theory of Cryptography (TCC)*, New York, NY, 2006, pp. 265–284.
- [2] C. Dwork, “Differential privacy,” in *Proc. Int. Conf. Automata, Languages and Programming (ICALP)*, Venice, Italy, 2006, pp. 1–12.
- [3] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor, “Our data, ourselves: privacy via distributed noise generation,” in *Proc. Annu. Int. Conf. Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, St. Petersburg, Russia, 2006, pp. 486–503.

- [4] A. Blum, K. Ligett, and A. Roth, “A learning theory approach to non-interactive database privacy,” in *Proc. Ann. ACM Symp. Theory of Computing (STOC)*, Victoria, Canada, 2008, pp. 609–618.
- [5] C. Dwork, M. Naor, O. Reingold, G. N. Rothblum, and S. Vadhan, “On the complexity of differentially private data release: efficient algorithms and hardness results,” in *Proc. Ann. ACM Symp. Theory of Computing (STOC)*, Bethesda, MD, 2009, pp. 381–390.
- [6] A. Roth and T. Roughgarden, “Interactive privacy via the median mechanism,” in *Proc. Ann. ACM Symp. Theory of Computing (STOC)*, Cambridge, MA, 2010, pp. 765–774.
- [7] C. Dwork, G. N. Rothblum, and S. Vadhan, “Boosting and differential privacy,” in *Proc. Ann. IEEE Symp. Found. Comput. Sci. (FOCS)*, Las Vegas, NV, 2010, pp. 51–60.
- [8] M. Hardt and G. N. Rothblum, “A multiplicative weights mechanism for privacy-preserving data analysis,” in *Proc. Ann. IEEE Symp. Found. Comput. Sci. (FOCS)*, Las Vegas, NV, 2010, pp. 61–70.
- [9] M. Hardt, K. Ligett, and F. McSherry, “A simple and practical algorithm for differentially private data release,” in *Advances Neural Information Processing Systems (NIPS)*, Lake Tahoe, NV, 2012, pp. 2348–2356.
- [10] A. Gupta, A. Roth, and J. Ullman, “Iterative constructions and private data release,” in *Proc. Conf. Theory of Cryptography (TCC)*, Sicily, Italy, 2012, pp. 339–356.
- [11] M. Gaboardi, E. J. G. Arias, J. Hsu, A. Roth, and Z. S. Wu, “Dual query: Practical private query release for high dimensional data,” in *Int. Conf. Machine Learning (ICML)*, Beijing, China, 2014.
- [12] O. Heffetz and K. Ligett, “Privacy and data-based research,” *J. Econ. Perspect.*, vol. 28, no. 2, pp. 75–98, 2014.
- [13] I. Dinur and K. Nissim, “Revealing information while preserving privacy,” in *Symp. Principles Database Systems (PODS)*, San Diego, CA, 2003, pp. 202–210.
- [14] A. B. Tsybakov, *Introduction to Nonparametric Estimation*. New York: Springer, 2009.
- [15] J. C. Duchi, M. I. Jordan, and M. J. Wainwright, “Local privacy and statistical minimax rates,” *arXiv:1302.3203 [math.ST]*, Feb. 2013.
- [16] —, “Local privacy and minimax bounds: Sharp rates for probability estimation,” *arXiv:1305.6000 [math.ST]*, May 2013.
- [17] M. Hardt and K. Talwar, “On the geometry of differential privacy,” in *Proc. Ann. ACM Symp. Theory of Computing (STOC)*, Cambridge, MA, 2010, pp. 705–714.
- [18] C. Li, M. Hay, V. Rastogi, G. Miklau, and A. McGregor, “Optimizing linear counting queries under differential privacy,” in *Symp. Principles Database Systems (PODS)*, Indianapolis, IN, 2010, pp. 123–134.
- [19] J. Ullman, “Answering $n^{2+o(1)}$ counting queries with differential privacy is hard,” in *Proc. Ann. ACM Symp. Theory of Computing (STOC)*, Palo Alto, CA, 2013, pp. 361–370.
- [20] F. McSherry and K. Talwar, “Mechanism design via differential privacy,” in *Proc. Ann. IEEE Symp. Found. Comput. Sci. (FOCS)*, Providence, RI, 2007, pp. 94–103.
- [21] K. B. Athreya and S. N. Lahiri, *Measure Theory and Probability Theory*. New York, NY: Springer, 2006.
- [22] “Netflix Prize,” <http://www.netflixprize.com>.
- [23] J. J. McAuley and J. Leskovec, “Learning to discover social circles in ego networks,” in *Advances Neural Information Processing Systems (NIPS)*, Lake Tahoe, NV, 2012, pp. 548–556.
- [24] J. Blocki, A. Blum, A. Datta, and O. Sheffet, “The johnson-lindenstrauss transform itself preserves differential privacy,” in *Proc. Ann. IEEE Symp. Found. Comput. Sci. (FOCS)*, New Brunswick, NJ, 2012, pp. 410–419.
- [25] A. Gupta, M. Hardt, A. Roth, and J. Ullman, “Privately releasing conjunctions and the statistical query barrier,” in *Proc. Ann. ACM Symp. Theory of Computing (STOC)*, San Jose, CA, 2011, pp. 803–812.

- [26] W. Wang, L. Ying, and J. Zhang, “On the relation between identifiability, differential privacy and mutual-information privacy,” in *Proc. Ann. Allerton Conf. Communication, Control and Computing*, Monticello, IL, 2014.
- [27] K. L. Chung, *A Course in Probability Theory*, 3rd ed. San Diego, CA: Academic Press, 2000.

APPENDIX A

PROOF OF LEMMA 2

Proof. We drop the subscript $Y \sim \mu_{\mathcal{E}(x)}$ from expectations for conciseness during the proof. We first prove that the estimator \hat{q}_φ^u is unbiased. Recall that $\{Y_i, i \in [n]\}$ follow the pmfs in (16). Then

$$\begin{aligned}
 \mathbb{E}[q_\varphi(Y)] &= \frac{1}{\sum_{i=1}^n c_i} \sum_{i=1}^n \mathbb{E}[\varphi_i(Y_i)] \\
 &= \frac{1}{\sum_{i=1}^n c_i} \sum_{i=1}^n \left(\frac{1}{g(\epsilon)} \varphi_i(x_i) + \frac{e^{-\epsilon}}{g(\epsilon)} \sum_{\substack{v \in \mathcal{D}: \\ v \neq x_i}} \varphi_i(v) \right) \\
 &= \frac{1 - e^{-\epsilon}}{g(\epsilon)} \frac{1}{\sum_{i=1}^n c_i} \sum_{i=1}^n \varphi_i(x_i) \\
 &\quad + \frac{e^{-\epsilon}}{g(\epsilon)} \frac{1}{\sum_{i=1}^n c_i} \sum_{i=1}^n \sum_{v \in \mathcal{D}} \varphi_i(v) \\
 &= \frac{1 - e^{-\epsilon}}{g(\epsilon)} q_\varphi(x) + \frac{e^{-\epsilon}}{g(\epsilon)} C_\varphi.
 \end{aligned}$$

Therefore

$$\mathbb{E}[\hat{q}_\varphi^u(Y)] = \mathbb{E}\left[\frac{g(\epsilon)}{1 - e^{-\epsilon}} q_\varphi(Y) - \frac{e^{-\epsilon}}{1 - e^{-\epsilon}} C_\varphi\right] = q_\varphi(x).$$

Next we prove the upper bound on the distortion of \hat{q}_φ^u . For any $x \in \mathcal{D}^n$,

$$\begin{aligned}
 &\hat{q}_\varphi^u(Y) - q_\varphi(x) \\
 &= \frac{g(\epsilon)}{1 - e^{-\epsilon}} \frac{1}{\sum_{i=1}^n c_i} \\
 &\quad \cdot \sum_{i=1}^n \left(\varphi_i(Y_i) - \frac{1 - e^{-\epsilon}}{g(\epsilon)} \varphi_i(x_i) - \frac{e^{-\epsilon}}{g(\epsilon)} \sum_{v \in \mathcal{D}} \varphi_i(v) \right).
 \end{aligned}$$

For any $i \in [n]$, let

$$Z_i = \varphi_i(Y_i) - \frac{1 - e^{-\epsilon}}{g(\epsilon)} \varphi_i(x_i) - \frac{e^{-\epsilon}}{g(\epsilon)} \sum_{v \in \mathcal{D}} \varphi_i(v).$$

Then for any $i \in [n]$, $\mathbb{E}[Z_i] = 0$. Recall that for any $v \in \mathcal{D}$, $a \leq \varphi_i(v) \leq b$, so $|Z_i| \leq b - a$. Since

Y_1, \dots, Y_n are independent, Z_1, \dots, Z_n are independent. Let $\bar{Z} = \frac{1}{n} \sum_{i=1}^n Z_i$. Then

$$\begin{aligned} & \mathbb{E}[|\hat{q}_\varphi^u(Y) - q_\varphi(x)|^2] \\ &= \left(\frac{g(\epsilon)}{1 - e^{-\epsilon}} \frac{n}{\sum_{i=1}^n c_i} \right)^2 \cdot \mathbb{E}[|\bar{Z}|^2] \\ &= \left(\frac{g(\epsilon)}{1 - e^{-\epsilon}} \frac{n}{\sum_{i=1}^n c_i} \right)^2 \cdot \left(\frac{1}{n^2} \sum_{i=1}^n \mathbb{E}[|Z_i|^2] \right) \\ &\leq \left(\frac{g(\epsilon)}{1 - e^{-\epsilon}} \right)^2 \frac{1}{c^2} \frac{(b-a)^2}{n}. \end{aligned}$$

Therefore

$$\sup_{x \in \mathcal{D}^n} \mathbb{E}[|\hat{q}_\varphi^u(Y) - q_\varphi(x)|^2] \leq \frac{(b-a)^2 (1 + (2^l - 1)e^{-\epsilon})^2}{c^2 (1 - e^{-\epsilon})^2} \frac{1}{n}.$$

□

APPENDIX B

PROOF OF LEMMA 3

Proof. For any $x, y \in \mathcal{D}^n$, since $q_\varphi(x) \in q_\varphi(\mathcal{D}^n)$, by the definition of the estimator \hat{q}_φ in (21),

$$|\hat{q}_\varphi^u(y) - \hat{q}_\varphi(y)| \leq |\hat{q}_\varphi^u(y) - q_\varphi(x)|.$$

Therefore

$$\begin{aligned} |\hat{q}_\varphi(y) - q_\varphi(x)| &\leq |\hat{q}_\varphi(y) - \hat{q}_\varphi^u(y)| + |\hat{q}_\varphi^u(y) - q_\varphi(x)| \\ &\leq 2|\hat{q}_\varphi^u(y) - q_\varphi(x)|, \end{aligned}$$

and

$$\mathbb{E}[|\hat{q}_\varphi(Y) - q_\varphi(x)|^2] \leq 4\mathbb{E}[|\hat{q}_\varphi^u(Y) - q_\varphi(x)|^2].$$

Then combining with (20) yields the upper bound. □

APPENDIX C

PROOF OF LEMMA 4

Proof. By Jensen's inequality,

$$\begin{aligned} & \mathbb{E}[\mathbb{E}[d(X, Z) \mid Y, Z] - d(X, Z)]^2 \\ &\geq (\mathbb{E}[\mathbb{E}[d(X, Z) \mid Y, Z] - d(X, Z)])^2. \end{aligned} \tag{37}$$

Let \tilde{X} be a random variable satisfying the following conditions: \tilde{X} is independent of Z ; \tilde{X} is independent of X given Y ; given Y , \tilde{X} and X are identically distributed, i.e., $p_{\tilde{X}|Y}(x | y) = p_{X|Y}(x | y)$ for any $x, y \in \mathcal{D}^n$ with $p_Y(y) \neq 0$. Due to the independence between Z and (X, Y, \tilde{X}) , we also have $p_{\tilde{X}|Y,Z}(x | y, z) = p_{X|Y,Z}(x | y, z)$ for any $x, y, z \in \mathcal{D}^n$ with $p_Y(y) \neq 0$. By this construction, for any $y, z \in \mathcal{D}^n$ with $p_Y(y) \neq 0$,

$$\mathbb{E}[d(X, Z) | Y = y, Z = z] = \mathbb{E}[d(\tilde{X}, Z) | Y = y, Z = z],$$

and

$$\begin{aligned} & \mathbb{E}[|\mathbb{E}[d(X, Z) | Y, Z] - d(X, Z)| | Y = y, Z = z] \\ &= \mathbb{E}[|\mathbb{E}[d(\tilde{X}, Z) | Y, Z] - d(\tilde{X}, Z)| | Y = y, Z = z], \end{aligned}$$

which further lead to

$$\begin{aligned} & \mathbb{E}[|\mathbb{E}[d(X, Z) | Y, Z] - d(X, Z)|] \\ &= \mathbb{E}\left[\mathbb{E}[|\mathbb{E}[d(X, Z) | Y, Z] - d(X, Z)| | Y, Z]\right] \\ &= \mathbb{E}\left[\mathbb{E}[|\mathbb{E}[d(\tilde{X}, Z) | Y, Z] - d(\tilde{X}, Z)| | Y, Z]\right] \\ &= \mathbb{E}[|\mathbb{E}[d(\tilde{X}, Z) | Y, Z] - d(\tilde{X}, Z)|]. \end{aligned}$$

Therefore

$$\begin{aligned} & 2\mathbb{E}[|\mathbb{E}[d(X, Z) | Y, Z] - d(X, Z)|] \\ &= \mathbb{E}[|\mathbb{E}[d(X, Z) | Y, Z] - d(X, Z)| \\ & \quad + |\mathbb{E}[d(\tilde{X}, Z) | Y, Z] - d(\tilde{X}, Z)|] \\ &\geq \mathbb{E}[|d(X, Z) - d(\tilde{X}, Z)| \\ & \quad + |\mathbb{E}[d(\tilde{X}, Z) | Y, Z] - \mathbb{E}[d(X, Z) | Y, Z]|] \\ &= \mathbb{E}[|d(X, Z) - d(\tilde{X}, Z)|]. \end{aligned}$$

Combing this with (37) gives

$$\begin{aligned} & \mathbb{E}[|\mathbb{E}[d(X, Z) | Y, Z] - d(X, Z)|^2] \\ &\geq \frac{1}{4}(\mathbb{E}[|d(X, Z) - d(\tilde{X}, Z)|])^2. \end{aligned} \tag{38}$$

Then it suffices to derive a lower bound on $\mathbb{E}[|d(X, Z) - d(\tilde{X}, Z)|]$.

Notice that the conditional pmf $p_{\tilde{X}|X}$ is ϵ -differentially private since for any neighboring $x, x' \in \mathcal{D}^n$ and any $\tilde{x} \in \mathcal{D}^n$,

$$p_{\tilde{X}|X}(\tilde{x} | x) = \sum_{y \in \mathcal{D}^n} p_{\tilde{X}|Y,X}(\tilde{x} | y, x) p_{Y|X}(y | x) \quad (39)$$

$$= \sum_{y \in \mathcal{D}^n} p_{\tilde{X}|Y,X}(\tilde{x} | y, x') p_{Y|X}(y | x) \quad (40)$$

$$\leq \sum_{y \in \mathcal{D}^n} p_{\tilde{X}|Y,X}(\tilde{x} | y, x') \cdot e^\epsilon p_{Y|X}(y | x') \quad (41)$$

$$= e^\epsilon p_{\tilde{X}|X}(\tilde{x} | x'), \quad (42)$$

where (40) follows from the conditional independence between \tilde{X} and X given Y , and (41) holds because $p_{Y|X}$ is ϵ -differentially private. Then by Theorem 1 in [26] (for our case, the ϵ_X in that theorem is 0),

$$\mathbb{E}[d(X, \tilde{X})] \geq \frac{n}{1 + \frac{e^\epsilon}{2^l - 1}}.$$

Let $\gamma = \frac{1}{2(1 + \frac{e^\epsilon}{2^l - 1})}$ and $s = \gamma n$. Since

$$\begin{aligned} \mathbb{E}[d(X, \tilde{X})] &\leq s \mathbb{P}\{d(X, \tilde{X}) < s\} + n \mathbb{P}\{d(X, \tilde{X}) \geq s\} \\ &\leq s + n \mathbb{P}\{d(X, \tilde{X}) \geq s\}, \end{aligned}$$

we have

$$\begin{aligned} \mathbb{P}\{d(X, \tilde{X}) \geq s\} &\geq \frac{1}{n} (\mathbb{E}[d(X, \tilde{X})] - s) \\ &\geq \frac{1}{n} \left(\frac{n}{1 + \frac{e^\epsilon}{2^l - 1}} - \frac{n}{2(1 + \frac{e^\epsilon}{2^l - 1})} \right) \\ &= \gamma, \end{aligned}$$

i.e.,

$$\mathbb{P}\{d(X, \tilde{X}) \geq \gamma n\} \geq \gamma. \quad (43)$$

We will consider those $x, \tilde{x} \in \mathcal{D}^n$ with $d(x, \tilde{x}) \geq \gamma n$ to obtain a lower bound on $\mathbb{E}[|d(X, Z) - d(\tilde{X}, Z)|]$.

Utilizing conditional expectation gives

$$\begin{aligned} &\mathbb{E}[|d(X, Z) - d(\tilde{X}, Z)|] \\ &= \mathbb{E}[\mathbb{E}[|d(X, Z) - d(\tilde{X}, Z)| | X, \tilde{X}]] \\ &\geq \sum_{\substack{x, \tilde{x}: \\ d(x, \tilde{x}) \geq \gamma n}} \mathbb{E}[|d(X, Z) - d(\tilde{X}, Z)| | X = x, \tilde{X} = \tilde{x}] p_{X, \tilde{X}}(x, \tilde{x}). \end{aligned} \quad (44)$$

Consider any $x, \tilde{x} \in \mathcal{D}^n$ with $d(x, \tilde{x}) \geq \gamma n$ and $p_{X, \tilde{X}}(x, \tilde{x}) \neq 0$. Since Z is independent of (X, \tilde{X}) ,

$$\begin{aligned} & \mathbb{E}[|d(X, Z) - d(\tilde{X}, Z)| \mid X = x, \tilde{X} = \tilde{x}] \\ &= \mathbb{E}[|d(x, Z) - d(\tilde{x}, Z)|]. \end{aligned} \quad (45)$$

Let

$$\Delta(x, \tilde{x}) = \{i \in [n] \mid x_i \neq \tilde{x}_i\}. \quad (46)$$

Then $|\Delta(x, \tilde{x})| \geq \gamma n$, and

$$\begin{aligned} |d(x, Z) - d(\tilde{x}, Z)| &= \left| \sum_{i=1}^n (\zeta(x_i, Z_i) - \zeta(\tilde{x}_i, Z_i)) \right| \\ &= \left| \sum_{i \in \Delta(x, \tilde{x})} (\zeta(x_i, Z_i) - \zeta(\tilde{x}_i, Z_i)) \right|. \end{aligned}$$

Let

$$U_i = \zeta(x_i, Z_i) - \zeta(\tilde{x}_i, Z_i). \quad (47)$$

Since Z is uniformly distributed over \mathcal{D}^n , the rows Z_1, Z_2, \dots, Z_n are i.i.d. with pmf $p_{Z_i}(z_i) = \frac{1}{2^l}$ for any $z_i \in \mathcal{D}$. For any $i \in \Delta(x, \tilde{x})$,

$$U_i = \begin{cases} 1 & \text{if } Z_i = \tilde{x}_i, \\ -1 & \text{if } Z_i = x_i, \\ 0 & \text{otherwise.} \end{cases} \quad (48)$$

Therefore $\{U_i, i \in \Delta(x, \tilde{x})\}$ are i.i.d. with pmf

$$p_{U_i}(u_i) = \begin{cases} \frac{1}{2^l} & u_i = 1, \\ \frac{1}{2^l} & u_i = -1, \\ 1 - \frac{1}{2^{l-1}} & u_i = 0. \end{cases} \quad (49)$$

Then $\mathbb{E}[U_i] = 0$. Denote

$$\sigma^2 = \mathbb{E}[|U_i|^2] = \frac{1}{2^{l-1}}, \quad \rho = \mathbb{E}[|U_i|^3] = \frac{1}{2^{l-1}}. \quad (50)$$

By the Berry–Esseen theorem [27, Theorem 7.4.1], there exists a universal constant C such that for any

t ,

$$\begin{aligned}
& \mathbb{P}\left\{\frac{1}{\sigma\sqrt{|\Delta(x, \tilde{x})|}} \sum_{i \in \Delta(x, \tilde{x})} U_i > \frac{t}{\sigma\sqrt{\gamma}}\right\} \\
& \geq 1 - \Phi\left(\frac{t}{\sigma\sqrt{\gamma}}\right) - \frac{C\rho}{\sigma^3\sqrt{|\Delta(x, \tilde{x})|}} \\
& \geq 1 - \Phi\left(\frac{t}{\sigma\sqrt{\gamma}}\right) - \frac{C\rho}{\sigma^3\sqrt{\gamma n}},
\end{aligned}$$

where the second inequality follows from $|\Delta(x, \tilde{x})| \geq \gamma n$. Therefore

$$\begin{aligned}
& \mathbb{P}\{|d(x, Z) - d(\tilde{x}, Z)| > t\sqrt{n}\} \\
& = \mathbb{P}\left\{\frac{1}{\sigma\sqrt{|\Delta(x, \tilde{x})|}} \sum_{i \in \Delta(x, \tilde{x})} U_i > \frac{t\sqrt{n}}{\sigma\sqrt{|\Delta(x, \tilde{x})|}}\right\} \\
& \geq \mathbb{P}\left\{\frac{1}{\sigma\sqrt{|\Delta(x, \tilde{x})|}} \sum_{i \in \Delta(x, \tilde{x})} U_i > \frac{t\sqrt{n}}{\sigma\sqrt{\gamma n}}\right\} \\
& \geq 1 - \Phi\left(\frac{t}{\sigma\sqrt{\gamma}}\right) - \frac{C\rho}{\sigma^3\sqrt{\gamma n}}.
\end{aligned}$$

Let $t = \sigma\sqrt{\gamma}$, then

$$\mathbb{P}\{|d(x, Z) - d(\tilde{x}, Z)| > \sigma\sqrt{\gamma n}\} \geq 1 - \Phi(1) - \frac{C\rho}{\sigma^3\sqrt{\gamma n}},$$

and further

$$\begin{aligned}
& \mathbb{E}[|d(x, Z) - d(\tilde{x}, Z)|] \\
& \geq \sigma\sqrt{\gamma n} \cdot \mathbb{P}\{|d(x, Z) - d(\tilde{x}, Z)| > \sigma\sqrt{\gamma n}\} \\
& \geq (1 - \Phi(1))\sigma\sqrt{\gamma n} - \frac{C\rho}{\sigma^3}.
\end{aligned} \tag{51}$$

Inserting this lower bound back to (45), (44) and combining the lower bound (43) yield

$$\begin{aligned}
& \mathbb{E}[|d(X, Z) - d(\tilde{X}, Z)|] \\
& \geq \sum_{\substack{x, \tilde{x}: \\ d(x, \tilde{x}) \geq \gamma n}} \left((1 - \Phi(1))\sigma\sqrt{\gamma n} - \frac{C\rho}{\sigma^3} \right) p_{X, \tilde{X}}(x, \tilde{x}) \\
& = \left((1 - \Phi(1))\sigma\sqrt{\gamma n} - \frac{C\rho}{\sigma^3} \right) \mathbb{P}\{d(X, \tilde{X}) \geq \gamma n\} \\
& \geq (1 - \Phi(1))\sigma\gamma^{\frac{3}{2}}\sqrt{n} - \frac{C\rho\gamma}{\sigma^3}.
\end{aligned}$$

Therefore, by (38),

$$\begin{aligned} & \mathbb{E}[|\mathbb{E}[d(X, Z) \mid Y, Z] - d(X, Z)|^2] \\ & \geq \frac{1}{4} \left((1 - \Phi(1)) \sigma \gamma^{\frac{3}{2}} \sqrt{n} - \frac{C \rho \gamma}{\sigma^3} \right)^2, \end{aligned}$$

which completes the proof. \square